arXiv:1407.4414v1 [stat.CO] 16 Jul 2014

Particle Filtering and Smoothing Using Windowed Rejection Sampling

J. N. Corcoran and D. Jennings
University of Colorado *

July 9, 2014

Abstract

"Particle methods" are sequential Monte Carlo algorithms, typically involving importance sampling, that are used to estimate and sample from joint and marginal densities from a collection of a, presumably increasing, number of random variables. In particular, a particle filter aims to estimate the current state X_n of a stochastic system that is not directly observable by estimating a posterior distribution $\pi(x_n|y_1, y_2, \ldots, y_n)$ where the $\{Y_n\}$ are observations related to the $\{X_n\}$ through some measurement model $\pi(y_n|x_n)$. A particle smoother aims to estimate a marginal distribution $\pi(x_i|y_1, y_2, \ldots, y_n)$ for $1 \leq i < n$. Particle methods are used extensively for hidden Markov models where $\{X_n\}$ is a Markov chain as well as for more general state space models.
Existing particle filtering algorithms are extremely fast and easy to implement. Although they suffer from issues of degeneracy and "sample impoverishment", steps can be taken to minimize these problems and overall they are excellent tools for inference. However, if one wishes to sample from a posterior distribution of interest, a particle filter is only able to produce dependent draws. Particle smoothing algorithms are complicated and far less robust, often requiring cumbersome post-processing, "forward-backward" recursions, and multiple passes through subroutines. In this paper we introduce an alternative algorithm for both filtering and smoothing that is based on rejection sampling "in windows" . We compare both speed and accuracy of the traditional particle filter and this "windowed rejection sampler" (WRS) for several examples and show that good estimates for smoothing distributions are obtained at no extra cost.

1 Introduction

Particle filters and smoothers are sequential Monte Carlo methods, typically importance sampling methods, that are often employed to sample from and provide estimates of the

*Postal Address: Department of Applied Mathematics, University of Colorado, Box 526 Boulder CO 80309-0526, USA; email: corcoran@colorado.edu, phone: 303-492-0685
[0]Keywords: particle filtering, rejection sampling, hidden Markov models
AMS Subject classification: 65C05, 65C60, 62G07

distribution of a set or subset of latent variables in a hidden Markov model given observations. They are constructed specifically to provide updated sampling and estimation when additional observations become available without reprocessing all observations.

Consider the hidden Markov model with underlying and unobserved states $\{X_n\}_{n=0}^{\infty}$, transition density $\pi(x_n|x_{n-1})$, and an initial distribution with density $\pi(x_0)$. (In this paper we will assume a continuous state space, though the sampling techniques described will apply in the discrete case as well.) Suppose that $\{Y_n\}_{n=1}^{\infty}$ represents a sequence of observable variables that are conditionally independent when the unobserved states are given and where each Y_n is related to the unobserved process through X_n and a "measurement model" density $\pi(y_n|x_n)$. Such a model is depicted in Figure 1(a).

Figure 1: Depiction of One and Two Layer Hidden Markov Models

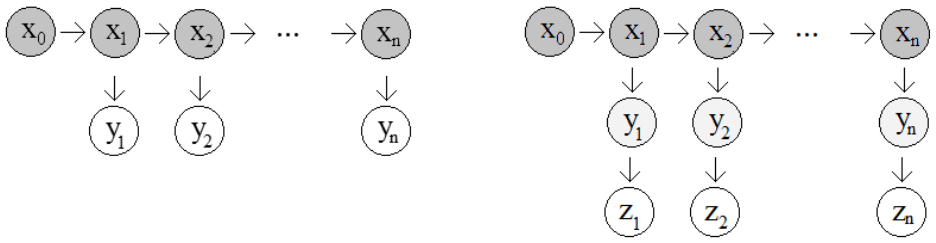

(a) (b)

Our goal is to sample from the density $\pi(x_{0:n}|y_{1:n})$, where $u_{i:j}$ denotes the vector $(u_i, u_{i+1}, \ldots, u_j)$ for $j \geq i$, sequentially in the sense that samples from $\pi(x_{0:n-1}|y_{1:n-1})$ will be used along with a new observation y_n to produce the desired points. The sampled points can then be used to approximate, for example, expectations of the form $\mathsf{E}[f(X_{0:n})|y_{1:n}]$ and marginal distributions $\pi(x_i|y_{1:n})$ for $i = 0, 1, \ldots, n$. The estimation of $\pi(x_n|y_{1:n})$ is known as the *filtering* problem and that of $\pi(x_i|y_{1:n})$ for $0 \leq i < n$ is known as the *smoothing* problem. In Section 4, we will see that our proposed algorithm, based on rejection sampling "in windows", vastly outperforms traditional particle methods for the smoothing problem and has some advantages for filtering as well.

The Markov and conditional independence assumptions allow us to write our "target density" for the model in Figure 1(a) in the recursive form

$$\pi(x_{0:n}|y_{1:n}) \propto \pi(x_0) \prod_{i=1}^{n} \pi(y_i|x_i)\pi(x_i|x_{i-1}) \tag{1a}$$

$$\propto \pi(y_n|x_n) \cdot \pi(x_n|x_{n-1}) \cdot \pi(x_{0:n-1}|y_{1:n-1}) \tag{1b}$$

so that draws from $\pi(x_{0:n}|y_{1:n})$ can be related to draws from $\pi(x_{0:n-1}|y_{1:n-1})$.

In Section 2, we review the sequential importance sampling and resampling methods, which make up the most commonly used particle filtering algorithm, that is typically used to

estimate and sample from (1a). We will then offer our new rejection sampling alternative in Section 3. In Section 4 we give simulation results, including direct comparisons to traditional particle filtering and an application to a second hidden layer Markov model.

2 Importance Sampling Based Sequential Methods

We now briefly describe existing importance sampling and resampling approaches to the problem of estimating and/or sampling sequentially from a target density

$$\pi(x_{1:n}) = h(x_{1:n})/Z_n, \tag{2}$$

where $Z_n = \int h(x_{1:n})\, dx_{1:n}$. A much more complete review is given by Doucet and Johansen [7]. Without loss of generality, $x_{1:n}$ can be replaced by $x_{0:n}$ and the target density can be a conditional density.

To establish notation, note that, if $X_{1:n}^{(1)}, X_{1:n}^{(2)}, \ldots, X_{1:n}^{(N)}$ represent N points ("particles") sampled from $\pi(x_{1:n})$, sequentially or otherwise, a simple unbiased estimator of the target density is given by the empirical density that puts weight $1/N$ on each of the points. We may write this succinctly as

$$\widehat{\pi}(x_{1:n}) = \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}[X_{1:n}^{(i)} = x_{1:n}] \tag{3}$$

where $\mathbb{1}[X_{1:n}^{(i)} = x_{1:n}]$ is the indicator function that takes on the value 1 when $x_{1:n} = X_{1:n}^{(i)}$, and zero otherwise.

We can then estimate, for example, $\mathsf{E}[f(X_{1:n})]$, for a given function f by

$$\widehat{\mathsf{E}}[f(X_{1:n})] = \frac{1}{N}\sum_{i=1}^{N} f(X_{1:n}^{(i)}). \tag{4}$$

It is easy to verify that (3) and (4) are unbiased for $\pi(x_{1:n})$ and $\mathsf{E}[f(X_{1:n})]$, respectively.

2.1 A Review of Non-Sequential Importance Sampling

When one can not sample directly from $\pi(x_{1:n}) \propto h(x_{1:n})$, importance sampling can be used to instead sample points from a more tractable density, and these points can be used to estimate the target density. Suppose that $q(x_{1:n})$ is another density, presumably tractable, with the same support as $\pi(x_{1:n})$.

We can write

$$\pi(x_{1:n}) = \frac{h(x_{1:n})}{Z_n} = \frac{w(x_{1:n})q(x_{1:n})}{Z_n} \tag{5}$$

where

$$w(x_{1:n}) := \frac{h(x_{1:n})}{q(x_{1:n})}. \tag{6}$$

Then we proceed to estimate $\pi(x_{1:n})$ and Z_n as follows.

- Sample $X_{1:n}^{(1)}, X_{1:n}^{(2)}, \ldots, X_{1:n}^{(N)} \overset{iid}{\sim} q(x_{1:n})$.[1]
- Estimate $q(x_{1:n})$ with

$$\widehat{q(x_{1:n})} = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}[X_{1:n}^{(i)} = x_{1:n}] \tag{7}$$

- Estimate $Z_n = \int w(x_{1:n}) q(x_{1:n})\, dx_{1:n} = \mathsf{E}[w(X_{1:n})]$, when $X_{1:n} \sim q(x_{1:n})$, with

$$\widehat{Z_n} = \frac{1}{N} \sum_{i=1}^{N} w(X_{1:n}^{(i)}). \tag{8}$$

- Estimate $\pi(x_{1:n})$ with

$$\begin{aligned} \widehat{\pi(x_{1:n})} &= \frac{w(x_{1:n})\widehat{q(x_{1:n})}}{\widehat{Z_n}} = \frac{\frac{1}{N}\sum_{i=1}^{N} w(X_{1:n}^{(i)}) \mathbb{1}[X_{1:n}^{(i)} = x_{1:n}]}{\widehat{Z_n}} \\ &= \sum_{i=1}^{N} W_n^{(i)}\, \mathbb{1}[X_{1:n}^{(i)} = x_{1:n}] \end{aligned} \tag{9}$$

  where

$$W_n^{(i)} = \frac{w(X_{1:n}^{(i)})}{\sum_{j=1}^{N} w(X_{1:n}^{(j)})}.$$

It is routine to show that (9) is an unbiased estimator of $\pi(x_{1:n})$ and that the optimal importance sampling density $q(x_{1:n})$, in terms of minimizing variance, is $q(x_{1:n}) \equiv \pi(x_{1:n})$.

2.2 Sequential Importance Sampling (SIS)

Sequential importance sampling (SIS) is importance sampling for $\pi(x_{1:n})$ in such a way where draws from $\pi(x_{1:n-1})$ are "extended" to n-dimensional points that are then reweighted to produce draws from $\pi(x_{1:n})$. To this end, the importance sampling density is chosen to have the form

$$q(x_{1:n}) = q(x_1) \prod_{i=2}^{n} q(x_i | x_{i-1})$$

so that it may be sampled from recursively.

The associated importance sample weights may then also be computed recursively since, for $n \geq 2$,

$$\begin{aligned} w(x_{1:n}) &= \frac{h(x_{1:n})}{q(x_{1:n})} = \frac{h(x_{1:n})}{q(x_{1:n-1}) q(x_n | x_{n-1})} \\ &= \frac{h(x_{1:n-1})}{q(x_{1:n-1})} \cdot \frac{h(x_{1:n})}{h(x_{1:n-1}) q(x_n | x_{n-1})} \\ &=: w(x_{1:n-1}) \cdot \alpha(x_{1:n}) \end{aligned} \tag{10}$$

[1] "iid" denotes "independent and identically distributed" values.

where $\alpha(x_{1:n})$ is an *incremental weight function* that is defined as

$$\alpha(x_{1:n}) := \frac{h(x_{1:n})}{h(x_{1:n-1})q(x_n|x_{n-1})}.$$

SIS Algorithm

At the first time step ($n = 1$),

- Sample $X_1^{(1)}, X_1^{(2)}, \ldots, X_1^{(N)} \overset{iid}{\sim} q(x_1)$.
- Compute weights $w(X_1^{(i)})$ for $i = 1, 2, \ldots, N$ using (6).

At times $n \geq 2$,

- Sample $X_n^{(1)}, X_n^{(2)}, \ldots, X_n^{(N)}$ independently with $X_n^{(i)} \sim q(x_n|X_{n-1}^{(i)})$.
- Compute weights $w(X_{1:n}^{(i)}) = w(X_{1:n-1}^{(i)}) \cdot \alpha(X_{1:n}^{(i)})$ for $i = 1, 2, \ldots, N$.

At any time n, one can estimate $\pi(x_{1:n})$ and Z_n using (9) and (8). One can also obtain approximate dependent draws from $\pi(x_{1:n})$ by sampling from (9). That is, by sampling from the set of values $\{X_{1:n}^{(1)}, X_{1:n}^{(2)}, \ldots, X_{1:n}^{(N)}\}$ using respective weights $\{W_n^{(1)}, W_n^{(2)}, \ldots, W_n^{(N)}\}$.

In practice, iteration of the SIS algorithm leads to a "degeneracy of weights" problem (see, for example, [2], [5], [6], and [9]) where the weights of all but one particle will approach zero, causing the method to break down and give meaningless results.

One way to avoid the issue of degenerate weights is to implement a resampling scheme at each time step.

SIS Algorithm With Resampling (SIR)

At the first time step ($n = 1$),

- Sample $X_1^{(1)}, X_1^{(2)}, \ldots, X_1^{(N)} \overset{iid}{\sim} q(x_1)$.
- Compute weights $w(X_1^{(i)})$ for $i = 1, 2, \ldots, N$.

- Compute the normalized weights

$$W_1^{(i)} = \frac{w(X_1^{(i)})}{\sum_{j=1}^N w(X_1^{(j)})}$$

for $i = 1, 2, \ldots, N$.

- Sample N points, $\widetilde{X}_1^{(1)}, \widetilde{X}_1^{(2)}, \ldots, \widetilde{X}_1^{(N)}$, with replacement, from the set $\{X_1^{(1)}, X_1^{(2)}, \ldots, X_1^{(N)}\}$ with respective probabilities $\{W_1^{(1)}, W_1^{(2)}, \ldots, W_1^{(N)}\}$. That is, sample $\widetilde{X}_1^{(1)}, \widetilde{X}_1^{(2)}, \ldots, \widetilde{X}_1^{(N)}$ from

$$\widehat{\pi(x_1)} = \sum_{i=1}^N W_1^{(i)} \mathbb{1}[X_1^{(i)} = x_1].$$

Note that $\widetilde{X}_1^{(1)}, \widetilde{X}_1^{(2)}, \ldots, \widetilde{X}_1^{(N)}$ are now equally weighted particles, each with weight $1/N$.

Assign weights $w(\widetilde{X}_1^{(i)}) = 1/N$ for $i = 1, 2, \ldots, N$.

At times $n \geq 2$,

- Sample $X_n^{(1)}, X_n^{(2)}, \ldots, X_n^{(N)}$ independently with $X_n^{(i)} \sim q(x_n | \widetilde{X}_{n-1}^{(i)})$.

- Extend each particle $\widetilde{X}_{1:n-1}^{(i)}$ to particles $(\widetilde{X}_{1:n-1}^{(i)}, X_n^{(i)})$.

- Compute associated weights $w(\widetilde{X}_{1:n-1}^{(i)}, X_n^{(i)}) := \alpha(\widetilde{X}_{1:n-1}^{(i)}, X_n^{(i)})$ for $i = 1, 2, \ldots, N$.
(This is consistent with SIS since the previous weights in (10) have been replaced by the constant $1/N$ and the current weights have yet to be normalized.)

- Compute the normalized weights

$$W_n^{(i)} = \frac{w(\widetilde{X}_{1:n-1}^{(i)}, X_n^{(i)})}{\sum_{j=1}^N w(\widetilde{X}_{1:n-1}^{(j)}, X_n^{(j)})}$$

for $i = 1, 2, \ldots, N$.

- Sample N n-dimensional points, $\widetilde{X}_{1:n}^{(1)}, \widetilde{X}_{1:n}^{(2)}, \ldots, \widetilde{X}_{1:n}^{(N)}$, with replacement, from the set $\{(\widetilde{X}_{1:n-1}^{(1)}, X_n^{(i)}), (\widetilde{X}_{1:n-1}^{(2)}, X_n^{(2)}), \ldots, (\widetilde{X}_{1:n-1}^{(N)}, X_n^{(N)})\}$ with respective probabilities $\{W_n^{(1)}, W_n^{(2)}, \ldots, W_n^{(N)}\}$. That is, sample $\widetilde{X}_{1:n}^{(1)}, \widetilde{X}_{1:n}^{(2)}, \ldots, \widetilde{X}_{1:n}^{(N)}$ from

$$\widehat{\pi(x_{1:n})} = \sum_{i=1}^N W_n^{(i)} \mathbb{1}[(\widetilde{X}_{1:n-1}^{(i)}, X_n^{(i)}) = x_{1:n}]. \tag{11}$$

Note that $\widetilde{X}_{1:n}^{(1)}, \widetilde{X}_{1:n}^{(2)}, \ldots, \widetilde{X}_{1:n}^{(N)}$ are now equally weighted particles, each with weight $1/N$.

Assign weights $w(\widetilde{X}_{1:n}^{(i)}) = 1/N$ for $i = 1, 2, \ldots, N$.

As with SIS, we may, at any time n, estimate $\pi(x_{1:n})$ using (11) or by using

$$\widehat{\pi(x_{1:n})} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}[\widetilde{X}_{1:n}^{(i)} = x_{1:n}]. \tag{12}$$

We may obtain approximate dependent draws from $\pi(x_{1:n})$ by sampling from (12). That is, by sampling uniformly, with replacement, from the set of values $\{\widetilde{X}_{1:n}^{(1)}, \widetilde{X}_{1:n}^{(2)}, \ldots, \widetilde{X}_{1:n}^{(N)}\}$.

An obvious issue with the SIR algorithm, known as the *sample impoverishment problem*, is a decreasing diversity of particles as those with higher weights will likely be drawn multiple times. Additionally, the resampling step may make the algorithm prohibitively slow and can lead to increased variance of estimates. One way to address the speed and variance is to only resample in the case of a small (user chosen threshold for) *effective sample size*

$$N_{eff} := \left[\sum_{i=1}^{N} (W_n^{(i)})^2\right]^{-1}. \tag{13}$$

A small effective sample size indicates higher variability of the weights which, in turn, indicates impending degeneracy. In this case, resampling would be especially prudent.

It is important to note that, when resampling, one should be originally sampling a very large number of points in order to minimize the effects of resampling repeated values from a discrete distribution. Additional details about SIS and SIR procedures can be found in Doucet *at al.* [5, 6] and in Andrieu and Doucet [1].

3 The Windowed Rejection Sampler

In order to draw from a possibly unnormalized density $\pi(x) \propto h(x)$ using rejection sampling, one must find a density $q(x)$ and a constant M such that

$$h(x) \leq Mq(x).$$

Then, one proceeds as in Algorithm 1.

Algorithm 1 Rejection Sampling from $\pi(x)$:

1. Generate $X \sim q$.
2. Generate $U \sim \text{Uniform}(0, 1)$.
3. If $U > \frac{h(X)}{Mq(X)}$, return to 1. Otherwise, if $U \leq \frac{h(X)}{Mq(X)}$, X is a draw from π.

Since $\pi(x)$ may be a high-dimensional and conditional density, we could, in theory, use

rejection sampling to draw exactly from

$$\begin{aligned}\pi(x_{0:n}|y_{1:n}) \propto h(x_{0:n}|y_{1:n}) &:= \pi(x_0)\prod_{i=1}^{n}\pi(y_i|x_i)\pi(x_i|x_{i-1}) \\ &= \underbrace{\left[\prod_{i=1}^{n}\pi(y_i|x_i^*)\right]}_{M}\cdot\underbrace{\pi(x_0)\prod_{i=1}^{n}\pi(x_i|x_{i-1})}_{q(x_{0:n})}\end{aligned}$$

where

$$x_i^* = \arg\max_{x_i}\pi(y_i|x_i).$$

For most applications of particle filtering in the literature, x_i^*, or at least an upper bound on $\pi(y_i|x_i)$, with respect to x_i, is easily attainable. However, for even small values of n, the acceptance probability, $\prod_{i=1}^{n}\pi(y_i|x_i)/\pi(y_i|x_i^*)$, can be prohibitively small and of course we are losing the benefit of sequential sampling. Thus, we propose rejection sampling "in windows".

Rejection Sampling in Windows

"Windowed Rejection Sampling" (WRS) is based on the idea that, depending on the covariance structure of the chain, at some point future observed y's will eventually cease to significantly affect earlier x's. For example, if $\pi(x_0|y_1,y_2,y_3,y_4) \approx \pi(x_0|y_1,y_2,y_3)$, we can write

$$\begin{aligned}\pi(x_0,x_1,x_2,x_3,x_4|y_1,y_2,y_3,y_4) &= \pi(x_1,x_2,x_3,x_4|x_0,y_1,y_2,y_3,y_4)\cdot\pi(x_0|y_1,y_2,y_3,y_4) \\ &\approx \pi(x_1,x_2,x_3,x_4|x_0,y_1,y_2,y_3,y_4)\cdot\pi(x_0|y_1,y_2,y_3),\end{aligned}$$

and so we can sample approximately from $\pi(x_{0:4}|y_{1:4})$ by first sampling X_0 from $\pi(x_0|y_{1:3})$ and then sampling $X_{1:4}$ from $\pi(x_1,x_2,x_3,x_4|x_0,y_1,y_2,y_3,y_4)$. Sampling from $\pi(x_0|y_{1:3})$ can be achieved by sampling $X_{0:3}$ from $\pi(x_{0:3}|y_{1:3})$ and considering only the x_0 values. In this example, we say that we are using rejection sampling in a "window of length 4".

More formally, the WRS algorithm is described in Algorithm 2 for a given window length w. We discuss the choice of this tuning parameter in Sections 4 and 5.

Algorithm 2 Windowed Rejection Sampling from $\pi(x_{0:n}|y_{1:n})$:

Find $x_i^* = \arg\max_{x_i} \pi(y_i|x_i)$, and set a window length parameter w.

1. Generate $X_{0:w-1} \sim \pi(x_{0:w-1}|y_{1:w-1})$ using rejection sampling with

$$q(x_{0:w-1}) = \pi(x_0) \prod_{i=1}^{w-1} \pi(x_i|x_{i-1})$$

and $M = \prod_{i=1}^{w} \pi(y_i|x_{i-1}^*)$. Set $m = 1$.

2. Generate $X_{m:m+w-1} \sim \pi(x_{m:m+w-1}|x_{m-1}, y_{1:m+w-1}) = \pi(x_{m:m+w-1}|x_{m-1}, y_{m:m+w-1})$ using rejection sampling with

$$q(x_{m:m+w-1}|x_{m-1}) = \prod_{i=m}^{m+w-1} \pi(x_i|x_{i-1})$$

and $M = \prod_{i=m}^{m+w-1} \pi(y_i|x_i^*)$.
Note that

$$\begin{aligned} \pi(x_{m:m+w-1}|x_{m-1}, y_{m:m+w-1}) &\propto h(x_{m:m+w-1}|x_{m-1}, y_{m:m+w-1}) \\ &:= \prod_{i=m}^{m+w-1} \pi(y_i|x_i)\pi(x_i|x_{i-1}) \leq M \cdot q(x_{m,n+w-1}|x_{m-1}). \end{aligned}$$

Set $m = m + 1$ and return to the beginning of Step 2. Continue until $m + w - 1 = n$.

Note that if it is not obtainable, $\pi(y_i|x_i^*)$ in expressions for M can be replaced by any upper bound.

4 Examples

4.1 One Dimensional Normals with One Hidden Layer

We begin, for the purpose of illustration, with a very explicit description of the algorithm with a window length of $w = 3$, for the simple model

$$X_0 \sim N(\mu_0, \sigma_0^2), \qquad X_{n+1} = aX_n + \sigma_X \varepsilon_{n+1}, \qquad Y_n = bX_n + \sigma_Y \nu_n$$

where $\{\varepsilon_n\} \overset{iid}{\sim} N(0,1)$ and $\{\nu_n\} \overset{iid}{\sim} N(0,1)$ are independent sequences. Assume that only the Y_n are observed.

Let $N(x; \mu, \sigma^2)$ denote the $N(\mu, \sigma^2)$ density.

We begin by using rejection sampling to produce N iid draws from $\pi(x_{0:2}|y_{1:2})$

$$\pi(x_{0:2}|y_{1:2}) \propto [\pi(y_1|x_1)\pi(y_2|x_2)] \cdot [\pi(x_2|x_1)\pi(x_1|x_0)\pi(x_0)]$$

using $q(x_{0:2}) = \pi(x_2|x_1)\pi(x_1|x_0)\pi(x_0) = N(x_2; ax_1, \sigma_X^2) \cdot N(x_1; ax_0, \sigma_X^2) \cdot N(x_0; \mu_0, \sigma_0^2)$.

It is easy to see that

$$x_i^* = \arg\max_{x_i} \pi(y_i|x_i) = \arg\max_{x_i} N(y_i; bx_i, \sigma_Y^2) = y_i/b \qquad \text{for } i = 1, 2, \ldots.$$

We repeatedly draw independent realizations $x_{0:2}$ of $X_{0:2} \sim q(x_{0:2})$ and u of $U \sim \text{Uniform}(0, 1)$ until the first time that

$$u \leq \frac{\pi(y_1|x_1)\pi(y_2|x_2)}{\pi(y_1|x_1^*)\pi(y_2|x_2^*)}.$$

Repeating this procedure N times, we collect only the values of x_0 and record them as $X_0^{(1)}, X_0^{(2)}, \ldots, X_0^{(N)}$.

Now, moving our window of length $w = 3$ to the right by 1, we use rejection sampling to produce N iid draws from $\pi(x_{1:3}|x_0, y_{1:3})$

$$\pi(x_{1:3}|x_0, y_{1:3}) \propto [\pi(y_1|x_1)\pi(y_2|x_2)\pi(y_3|x_3)] \cdot [\pi(x_3|x_2)\pi(x_2|x_1)\pi(x_1|x_0)]$$

using

$$q(x_{1:3}|x_0) = \pi(x_3|x_2)\pi(x_2|x_1)\pi(x_1|x_0) = N(x_3; ax_2, \sigma_X^2) \cdot N(x_2; ax_1, \sigma_X^2) \cdot N(x_1; ax_0, \sigma_X^2).$$

That is, for each $i = 1, 2, \ldots, N$, we repeatedly draw independent realizations $x_{1:3}$ of $X_{1:3} \sim q(x_{1:3}|X_0^{(i)})$ and u of $U \sim \text{Uniform}(0, 1)$ until the first time that

$$u \leq \frac{\pi(y_1|x_1)\pi(y_2|x_2)\pi(y_3|x_3)}{\pi(y_1|x_1^*)\pi(y_2|x_2^*)\pi(y_3|x_3^*)},$$

collecting the resulting value of x_1 and recording it as $X_1^{(i)}$.

Move the window of length $w = 3$ to the right by 1 to produce N draws from $\pi(x_{2:4}|x_1, y_{1:4}) = \pi(x_{2:4}|x_1, y_{2:4})$, and retain the resulting values for $X_2^{(i)}$ for $i = 1, 2, \ldots, N$.

Continuing in this manner, we generated $N = 100,000$ independent values for $X_{0:11}$ for this model using parameters $\mu_0 = 3.0$, $\sigma_0 = 2.0$, $a = 0.9$, $b = 1.2$, $\sigma_X = 3.0$, and $\sigma_Y = 2.3$. (We looked first at such a large N in order to really see the resulting distribution without being concerned about sampling variability.) For comparison, we produced N independent draws directly from the 11-dimensional distribution $\pi(x_{0:10}|y_{1:10})$ using 11-dimensional (non-windowed) rejection sampling. (As expected, the acceptance probabilities are quite small in this case and such direct simulation from the target density is not reasonably achieved for much larger n.) Finally, we produced N dependent draws approximately from $\pi(x_{0:10}|y_{1:10})$ using the SIR algorithm. Figure 2 shows the average marginal values for X_0 through X_{10} for each algorithm. Only the WRS results change between graphs, showing the anticipated increasing accuracy of the WRS algorithm as the window length increases. For the given model and parameters, a window length of $w = 3$ appears to be sufficient. Indeed, as the perfect high-dimensional rejection sampling algorithm is a special case of the WRS algorithm for fixed n and maximal window length, it is easy to code the WRS algorithm

once and run it first to get perfect draws from $\pi(x_{0:n}|y_{1:n})$ for the highest feasible n and then to run it with lower w, gradually increasing w until sample statistics such as those shown in Figure 2 reach the desired accuracy.

For this example, the WRS algorithm produced 100,000 iid draws from an approximation to $\pi(x_{0:11}|y_{1:10})$ whereas the SIR algorithm (resampling every time) produced dependent draws with only roughly 50% unique 11-dimensional values and only about 10% unique values for X_0 marginally. Figure 3 shows marginal distributions for X_7 produced by the WRS and SIR algorithms. The overlayed curves are the target normal densities that can be computed analytically for this simple illustrative model. The time to run the WRS and SIR algorithms were comparable. Coded in C++, the WRS algorithm completed in less than 2 seconds on a standard[2] laptop while the SIR algorithm completed in 8-10 seconds. The SIR algorithm would obviously speed up if resampling was not done at every step and would likely be the faster algorithm if both were programmed in R where resampling is efficient and accept/reject looping is inefficient.

When comparing the algorithms for larger n, it becomes crucial for efficiency to not resample at every time step of the SIR algorithm. As per the discussion at the end of Section 2, we resample only when the effective sample size N_{eff}, given by (13) falls below the a threshold N_{thres}. Following a suggestion in [6], we choose $N_{thres} = N/3$. Due to the Markov nature of the model, we are able to continue forward with the window length $w = 3$ for increasing n. When $n = 1,000$, the estimated filtering distribution $\pi(x_{1000}|y_{1:1000})$ is shown in Figure 4 for both algorithms. The SIR algorithm resulted in approximately 52% distinct values. An estimated smoothing distribution, namely an estimate of $\pi(x_{300}|y_{1:1000})$ is shown in Figure 5. It is well known (see [7]) that the SIR algorithm will produce severely impoverished samples for $\pi(x_i|y_{1:n})$ and i small relative to n, and in this example we are left with only 0.07% distinct values. Suggested approaches to mitigate degeneracy in SIR are discussed in [7], but the WRS algorithm performs well without modification.

4.2 Stochastic Volatility Model

We now consider the *stochastic volatility* model, often used in finance and also considered in [7], where

$$X_0 \sim N(0, \sigma^2/(1-\alpha^2)), \qquad X_{n+1} = \alpha X_n + \sigma\varepsilon_{n+1}, \qquad Y_n = \beta e^{X_n/2}\nu_n$$

where $\{\varepsilon_n\} \overset{iid}{\sim} N(0,1)$ and $\{\nu_n\} \overset{iid}{\sim} N(0,1)$ are independent sequences. Again we assume that only the Y_n are observed.

Following the same procedure in Section 4.1, using a general window length w, we need to define $q(x_{0:w-1})$, x_i^* and $q(x_{j:(j+w-1)}|x_{j-1})$. It is easy to show that

$$x_i^* = \arg\max_{x_i} \pi(y_i|x_i) = \arg\max_{x_i} N(y_i; 0, \beta^2 e^{x_i}) = \ln(y_i^2/\beta^2) \qquad \text{for } i = 1, 2, \ldots.$$

Taking

$$q(x_0) = \pi(x_0) = N(x_0; 0, \sigma^2/(1-\alpha^2)) \text{ and } q(x_i|x_{i-1}) = \pi(x_{i|x_{i-1}}) = N(x_i; \alpha x_{i-1}, \sigma^2),$$

[2]We wish to convey relative speed between algorithms in lieu of particular machine specifications.

we have

$$q(x_{0:w-1}) = q(x_0) \prod_{i=1}^{w-1} q(x_i|x_{i-1}) \tag{14}$$

and

$$q(x_{j:(j+w-1)}|x_{j-1}) = \prod_{i=j}^{j+w-1} q(x_i|x_{i-1}). \tag{15}$$

We ran Step 1 of the WRS algorithm (Algorithm 2) with $w = 11$ in order to get $N = 100,000$ perfect draws via rejection sampling from $\pi(x_{0:10}|y_{1:10})$. ($n = 10$ was chosen as the maximal value after which rejection sampling from $\pi(x_{0:n}|y_{1:n})$ became prohibitively slow.) We then ran the complete WRS algorithm with $n = 10$ and increasing w, starting from $w = 1$, until the marginal sample means from WRS for $X_i|y_{1:10}$ produced rough (non-simultaneous) two standard deviation confidence intervals which contained the perfect draw estimates of the means. That is, if $\hat{\mu}_i$ denotes the sample mean of the 100,000 values of X_i resulting from rejection sampling from $\pi(x_{0:10}|y_{1:10})$ and $\overline{X}_i$ and S_i^2 denote the sample means and variances of the marginal distributions produced by WRS, w was increased until $\hat{\mu}_i \in \overline{X}_i \pm 2S_i/\sqrt{N}$ for all $i \in \{0, 1, \ldots, 10\}$. This resulted in a window length of 5. Figure 6 shows the sample means for the marginal draws with WRS and $w = 5$ aligning with those from both perfect and SIR draws. Figure 7 shows the proportion of surviving distinct values for each marginal distribution using SIR, which is in contrast to the WRS algorithm giving 100% distinct values for all marginals. Figure 8 shows histograms of the marginal distribution of $X_5|y_{1:10}$ obtained from the draws from $\pi(x_{0:10}|y_{1:10})$ using 11-dimensional rejection sampling, the SIR algorithm, and the WRS algorithm with $w = 5$. The overlayed curve in each case is the density estimated by 100,000 perfect draws of X_5 marginalized from 11-dimensional rejection sampling. The SIR algorithm results stand out as the roughest estimate, though, again, it is well known [7] that this "straight" application of the SIR algorithm does not produce good smoothing estimates and that additional steps should be taken to improve it. The WRS algorithm, on the other hand gives a nice estimate of the smoothing distribution straight away with a speed that is comparable to the initial pass of the SIR algorithm. (By "comparable" we mean that, for the examples in this paper as well as others we tried, the WRS algorithm was either faster than the SIR algorithm or slower but still finishing only a few seconds behind. In the worst case, the WRS algorithm was about one minute behind the SIR algorithm. No particular effort was made to fully optimize either code.)

4.3 A "Highly Nonlinear" Model

In this section, we apply the WRS algorithm the "highly nonlinear" model considered in several papers including [3], [4], [6], and [8].

$$\begin{aligned} X_0 &\sim N(\mu, \sigma^2) \\ X_{n+1} &= 0.5X_n + \tfrac{25X_n}{1+X_n^2} + 8\cos(1.2n) + \varepsilon_{n+1} \\ Y_n &= 0.05X_n^2 + \nu_n \end{aligned}$$

for independent $\varepsilon_n \overset{iid}{\sim} N(0, \sigma_X^2)$ and $\nu_n \overset{iid}{\sim} N(0, \sigma_Y^2)$. The parameters were chosen as $\mu = 0$, $\sigma^2 = 5$, and $\sigma_X^2 = \sigma_Y^2 = 10$.

Taking

$$q(x_0) = \pi(x_0) = N(x_0; \mu, \sigma^2) \text{ and } q(x_i|x_{i-1}) = \pi(x_i|x_{i-1}) = N(x_i; , m(x_{i-1}), \sigma_X^2),$$

with $m(x_{i-1}) = 0.5x_{i-1} + 25x_{i-1}/(1 + x_{i-1}^2) + 8\cos[1.2(i-1)]$, we define $q(x_{0:w-1})$ and $q(x_{j:(j+w-1)}|x_{j-1})$ using (14) and (15).

It is easy to show that

$$x_i^* = \arg\max_{x_i} \pi(y_i|x_i) = \begin{cases} 0 & , \text{ if } y_i < 0 \\ \pm\sqrt{y_i/0.05} & , \text{ if } y_i \geq 0. \end{cases}$$

An analysis of marginal means such as those depicted in Figures 2 and 6 reveal that a window length of 4 is sufficient for this example. Figure 9 shows the proportion of surviving distinct values for each marginal distribution using SIR, which, again, is in contrast to the WRS algorithm giving 100% distinct values for all marginals. Figures 10 show histograms of $X_5|y_{1:10}$ produced by the WRS and SIR algorithms along with superimposed densities estimated by 11-dimensional rejection sampling. As expected, the WRS results are smoother. In addition, the 100,000 11-dimensional WRS draws are distinct and were produced in about 25 seconds, whereas the SIR results took closer to 1 minute. If the SIR algorithm, again slowed by resampling, had been coded in R, it undoubtedly would have been faster.

4.4 Two Layer Hidden Markov Model

In [1], the authors discuss a Rao-Blackwellized particle filtering technique, based on the SIR algorithm, for estimating $\pi(x_{0:n}|z_{1:n})$ in a two-layer hidden Markov model as depicted in Figure 1(b). It involves using SIR to first estimate $\pi(y_{1:n}|z_{1:n})$ and then estimating $\pi(x_{0:n}|z_{1:n})$ using the relationship

$$\pi(x_{0:n}|z_{1:n}) = \int \pi(x_{0:n}|y_{1:n})\pi(y_{1:n}|z_{1:n})$$

and draws from $\widehat{\pi}(y_{1:n}|z_{1:n})$. The WRS algorithm can be used here in place of the SIR algorithm.

In this example, however, we consider a particularly simple two-layer model from [1] which can be dealt with more directly. The "dynamic tobit model" is described as follows.

$$X_0 \sim N\left(0, \frac{\sigma_X^2}{1-\phi^2}\right)$$

$$X_{n+1} = \phi X_n + \sigma_X \varepsilon_{n+1}, \qquad Y_n = X_n + \sigma_Y \nu_n, \qquad Z_n = \max(0, Y_n)$$

where $\{\varepsilon_n\} \overset{iid}{\sim} N(0,1)$ and $\{\nu_n\} \overset{iid}{\sim} N(0,1)$ are independent sequences. Assume that only the Z_n are observed. The parameters were chosen as $\sigma_X^2 = 0.05$, $\sigma_Y^2 = 0.30$, and $\phi = 0.99$, as in [1].

Consider the case first when $n = 1$. If $z_1 > 0$, we wish to draw from

$$\pi(x_0, x_1|z_1) = \pi(x_0, x_1|y_1) \propto h(x_0, x_1|y_1) = \pi(y_1|x_1)\pi(x_1|x_0)\pi(x_0)$$

where $y_1 = z_1$.

To implement the SIR algorithm, we take $q(x_0, x_1) := \pi(x_1|x_0)\pi(x_0)$ and use the weight $w(x_{0:1}) := \pi(y_1|x_1)$. To implement rejection sampling, we use the fact that

$$h(x_0, x_1|y_1) \leq M \cdot q(x_0, x_1)$$

where $M = 1\Big/\sqrt{2\pi\sigma_Y^2}$ and, again, $q(x_0, x_1) := \pi(x_1|x_0)\pi(x_0)$.

If $z_1 = 0$, we wish to draw from

$$\pi(x_0, x_1|z_1) = \pi(x_0, x_1|y_1 < 0) \propto P_\pi(y_1 < 0|x_1)\pi(x_1|x_0)\pi(x_0)$$

where $P_\pi(y_1 < 0|x_1) = \int_{-\infty}^{0} \pi(y_1|x_1)\, dy_1$. We then have the unnormalized target density

$$h(x_0, x_1|z_1) = \Phi(-x_1/\sigma_Y) \cdot \pi(x_1|x_0)\pi(x_0) \tag{16}$$

where $\Phi(\cdot)$ denotes the cumulative distribution function for the standard normal distribution.

Taking $q(x_0, x_1) := \pi(x_1|x_0)\pi(x_0)$, we can implement the SIR algorithm to draw from (16) using $w(x_{0:1}) = \Phi(-x_1/\sigma_Y)$ and we can implement rejection sampling using the fact that $\Phi(-x_1/\sigma_Y) \leq M := 1$.

For general n, we use $q(x_n|x_{n-1}) = \pi(x_n|x_{n-1})$ for both algorithms. The incremental weight for the SIR algorithm is

$$\alpha(x_{1:n}) = \begin{cases} \pi_{Y|X}(z_n|x_n) & , \text{ if } z_n > 0 \\ \\ \Phi(-x_n/\sigma_Y) & , \text{ if } z_n = 0. \end{cases}$$

Here, $\pi_{Y|X}(\cdot|\cdot)$ is the conditional density for Y_n given X_n, previously identifiable through arguments alone.

For the WRS algorithm with window length w, we define M in Step 1 of Algorithm 2 as

$$M = \prod_{i=1}^{w} \left\{\pi_{Y|X}(z_i|x_i)\mathbb{1}[z_i > 0] + \Phi(-x_i/\sigma_Y)\mathbb{1}[z_i = 0]\right\},$$

and M in Step 2 as

$$M = \prod_{i=m}^{m+w-1} \left\{\pi_{Y|X}(z_i|x_i)\mathbb{1}[z_i > 0] + \Phi(-x_i/\sigma_Y)\mathbb{1}[z_i = 0]\right\}.$$

Since the underlying hidden Markov model is an AR(1) model with autocorrelation function $\rho(h) = \phi^h$, it is not surprising that, for $\phi = 0.99$, we would need w to be large for the WRS

algorithm. As in previous examples, we produced 100,000 independent draws directly from the 11-dimensional distribution $\pi(x_{0:10}|z_{1:10})$ using 11-dimensional, non-windowed, rejection sampling and then used windowed rejection sampling with increasing values of w until we matched marginal means. Figure 11 shows the matching means, as well as the matching means from the SIR algorithm. The turning point in this graph of means corresponds to the fact that z_3 was the only data point in $z_{1:10}$ that was positive. Both the WRS and SIR algorithms were quite slow for this example (on the order of 45 minutes) but were still of comparable speeds.

In Figure 12 we see that the marginal means for the SIR and WRS algorithms continue to coincide even for large n. Due to the Markov nature of the model, the window length determined with smaller n can still be used.

5 Conclusions

The WRS algorithm produces independent and identically distributed draws from an approximation to $\pi(x_{0:n}|y_{1:n})$ in a semi-sequential manner. That is, for each new observed y_n, the algorithm produces not only values for x_n but reproduces values for some previous x_i within a certain window length. At first glance, this redundancy, combined with the use of rejection sampling which has the potential to loop for some time, makes the WRS algorithm seem impractical. Indeed the WRS algorithm would be very inefficient if it were programmed in R, however in other languages the WRS and SIR algorithms ran at comparable speeds for all of the examples we tried including the benchmark examples from the SIR literature that were included in this paper. Ideally, one would compare speeds between the WRS algorithm programmed in a looping friendly language such as FORTRAN and SIR in a resampling friendly language such as R. When making speed comparisons in this paper, it was not our intentions to "split hairs" at this level but only to assert that the WRS can compete and can do so with the added benefit of full samples of distinct iid values.

As for the smoothing problem, our comparison with SIR was also not entirely fair, as it is well known (See, for example, [7].) that one should take additional steps such as those involving forward-backward recursions or "backward information" filters. However, the point to be made is that the WRS algorithm produces smoothing distributions without any additional effort.

The WRS algorithm is not a good alternative to traditional particle filters for non-Markovian models or in cases where the required minimal window length is so large as to result in inefficient rejection sampling.

References

[1] C. Andrieu and A. Doucet. Particle filtering for partially observed Gaussian state space models. *Journal of the Royal Statistical Society, Ser. B*, 64(4):827–836, 2002.

[2] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society, Ser. B*, 72(2):1–33, 2010.

[3] M.S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002.

[4] W. Chen, B.R. Bakshi, P.K. Goel, and S. Ungarala. Bayesian estimation via sequential monte carlo sampling: unconstrained nonlinear dynamic systems. *Ind. Eng. Chem. Res.*, 43(14):4012–4025, 2004.

[5] A. Doucet, J.F.G. de Freitas, and N.J. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer, New York, 2001.

[6] A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte Carlo methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, 2000.

[7] A. Doucet and A.M. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. In *Handbook of Nonlinear Filtering*. 2011. 656-704.

[8] N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *Radar and Signal Processing, IEE Proceedings F*, 140(2):107–113, 1993.

[9] A.J. Haug. *Bayesian Estimation and Tracking*. Wiley, New Jersey, 2012.

Figure 2: Section 4.1 Example: Means for marginal components in $100,000$ draws from $\pi(x_{0:10}|y_{1:10})$ using perfect 11-dimensional rejection sampling, the SIR algorithm, and the WRS algorithm with various window lengths.

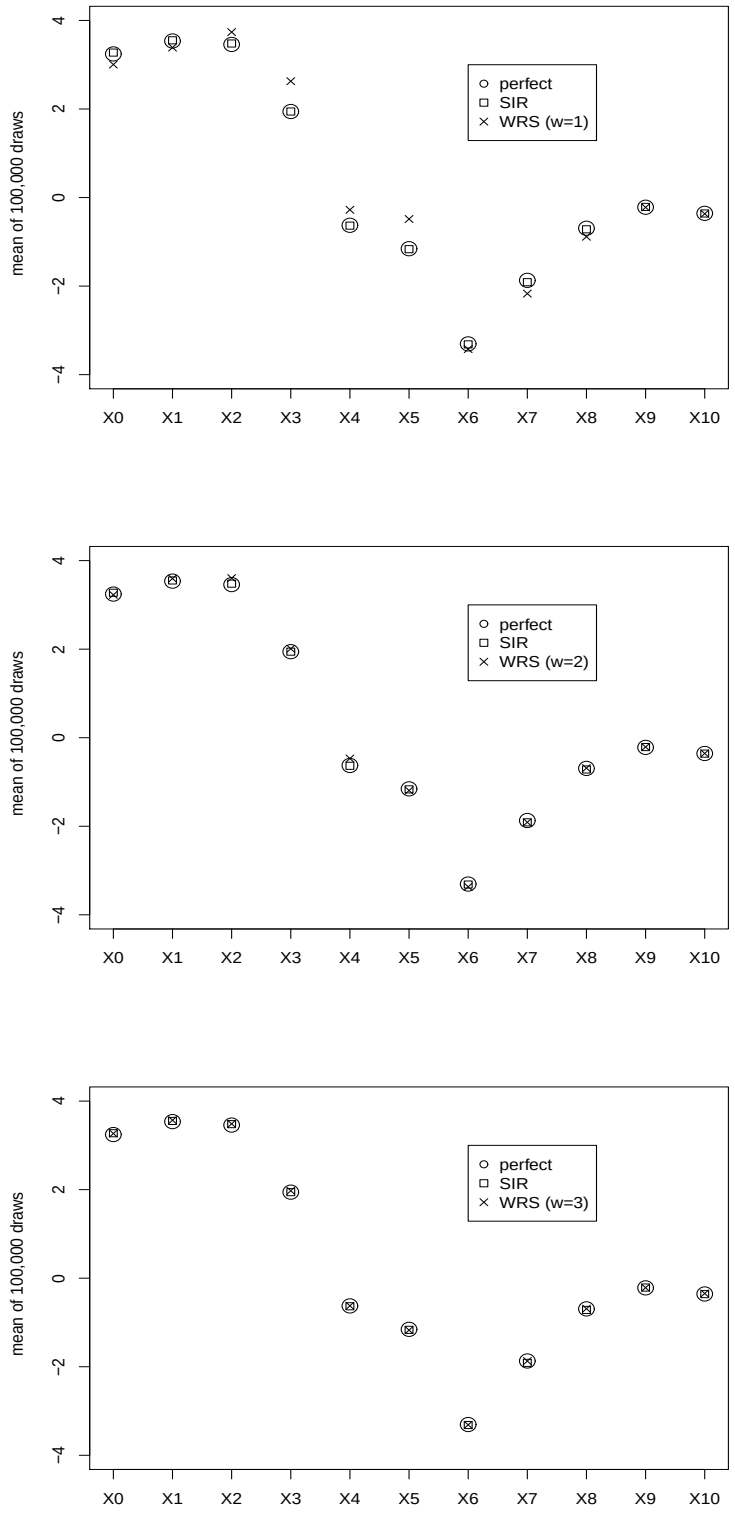

Figure 3: Section 4.1 Example: Histogram of values of X_7 from $100,000$ draws from $\pi(x_{0:11}|y_{1:10})$ using WRS and SIR algorithms

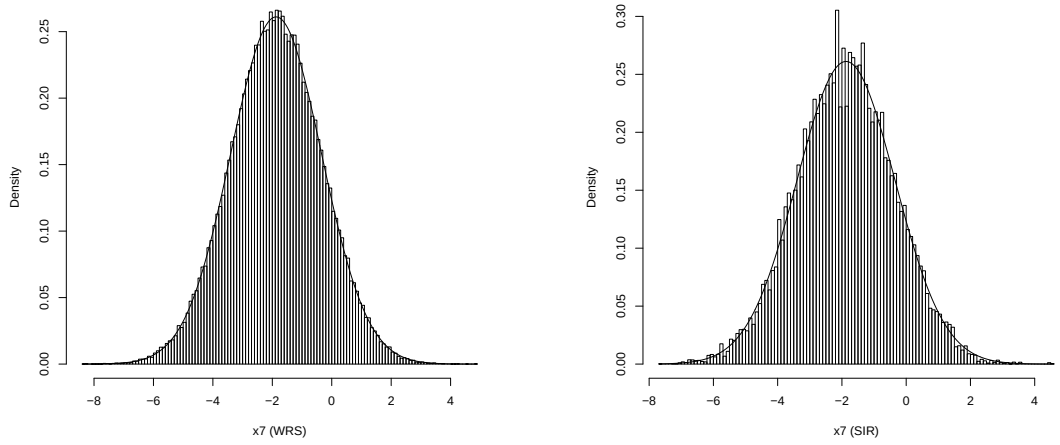

(a) WRS ($w = 3$)

(b) SIR ($\approx 21.3\%$ distinct)

Figure 4: Section 4.1 Example: Histogram of values of X_{1000} from $100,000$ draws from $\pi(x_{0:1000}|y_{1:10})$ using WRS and SIR algorithms

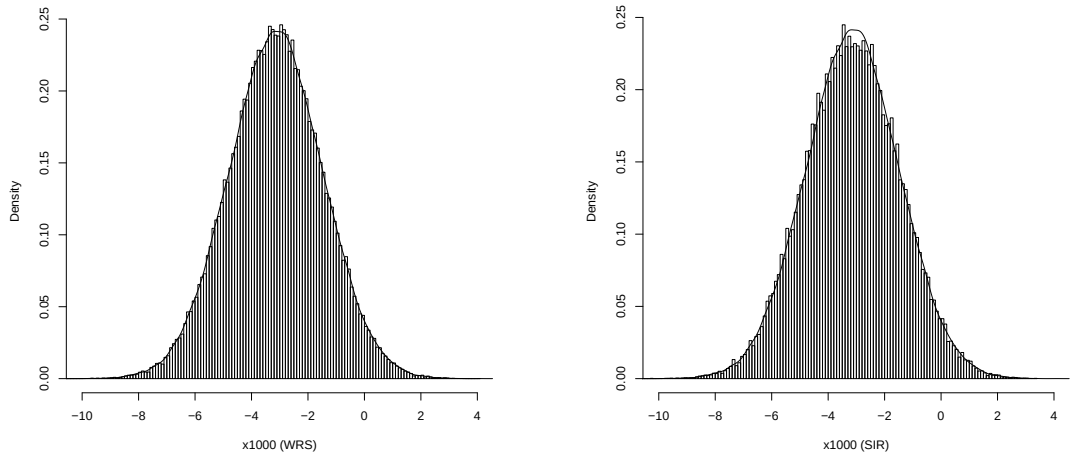

(a) WRS ($w = 3$)

(b) SIR ($\approx 52\%$ distinct)

Figure 5: Section 4.1 Example: Histogram of values of X_{300} from $100,000$ draws from $\pi(x_{0:1000}|y_{1:10})$ using WRS and SIR algorithms

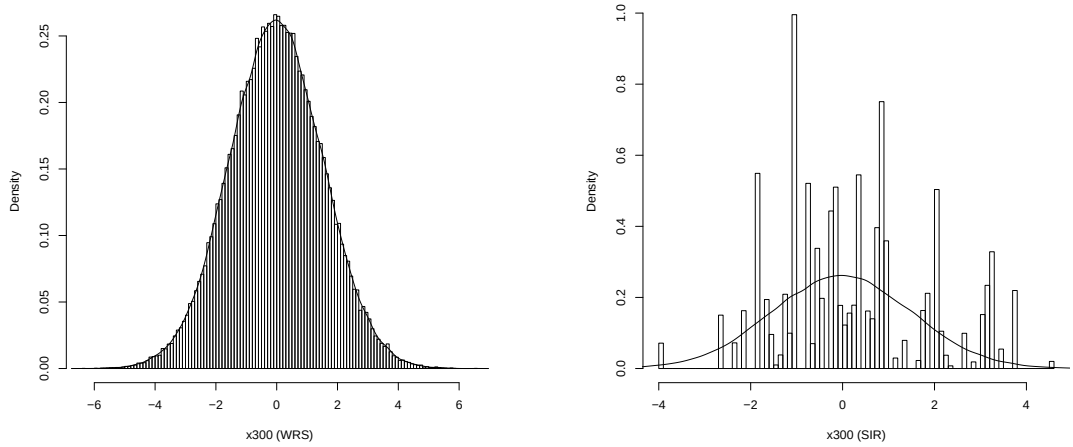

(a) WRS ($w = 3$)

(b) SIR (0.07% distinct)

Figure 6: Section 4.2 Example: Means for marginal components in $100,000$ draws from $\pi(x_{0:10}|y_{1:10})$ using perfect 11-dimensional rejection sampling, the SIR algorithm, and the WRS algorithm with $w = 5$.

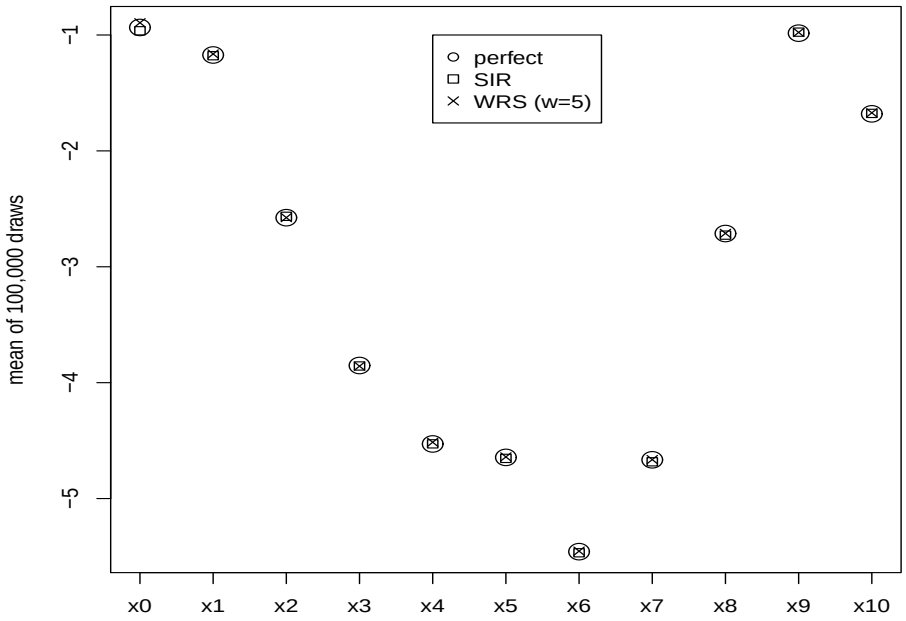

Figure 7: Section 4.2 Example: Proportion of surviving distinct marginal values using SIR

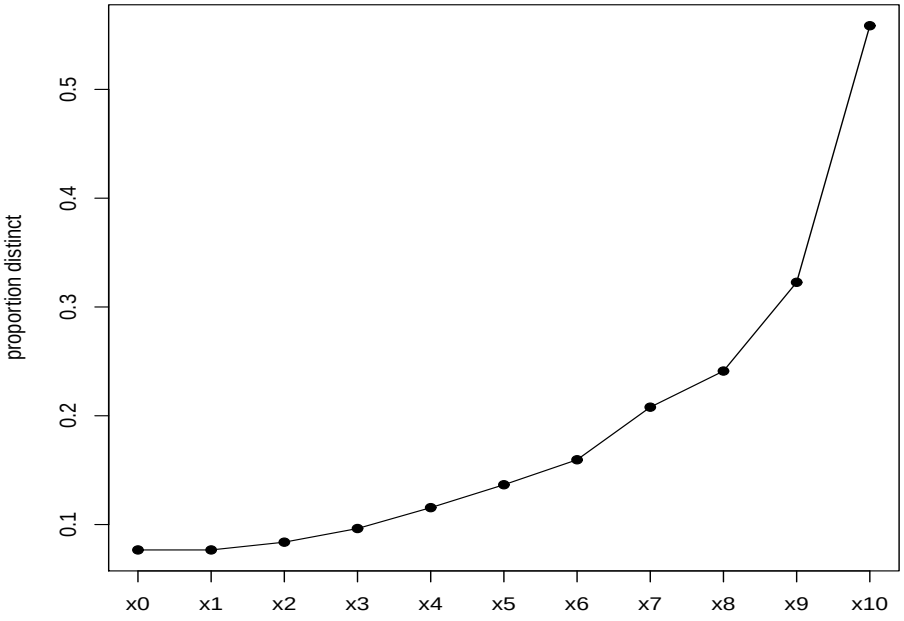

Figure 8: Section 4.2 Example: 100,000 draws from $X_5|y_{0:10}$ using WRS and SIR algorithms

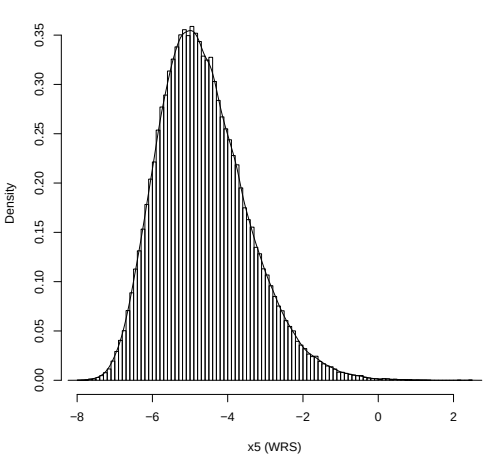

(a) WRS ($w = 5$)

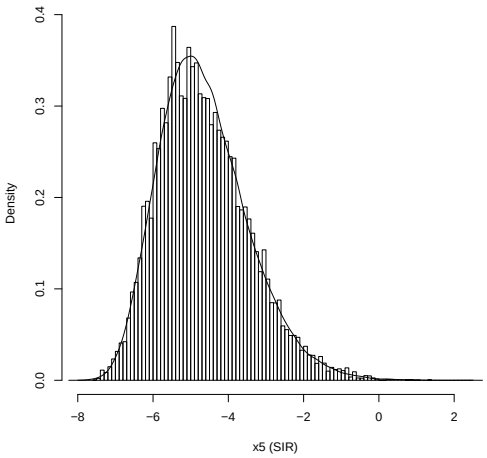

(b) SIR (13.6% distinct)

Figure 9: Section 4.3 Example: Proportion of Surviving Distinct Marginal Values using SIR

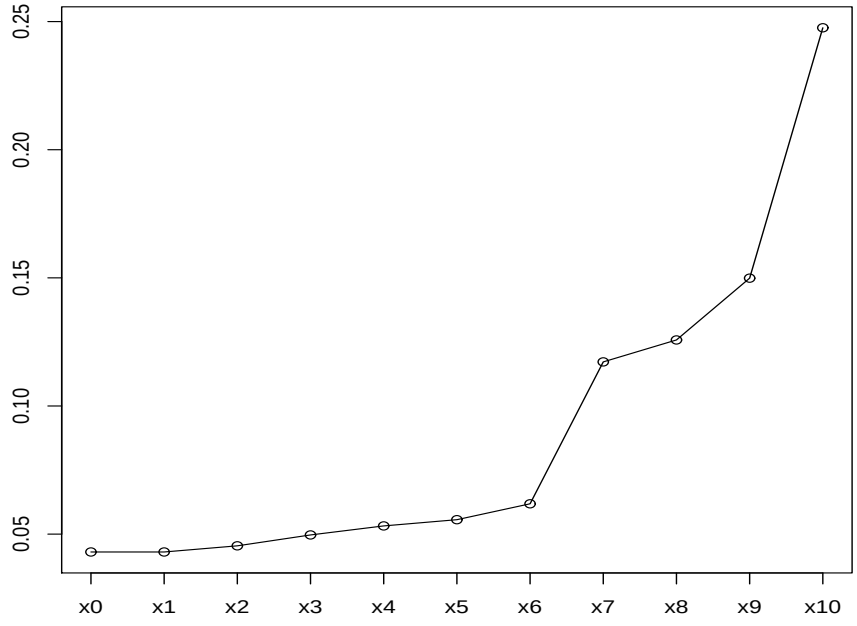

Figure 10: Section 4.3 Example: 100,000 draws from $X_5|y_{1:10}$ using WRS and SIR algorithms

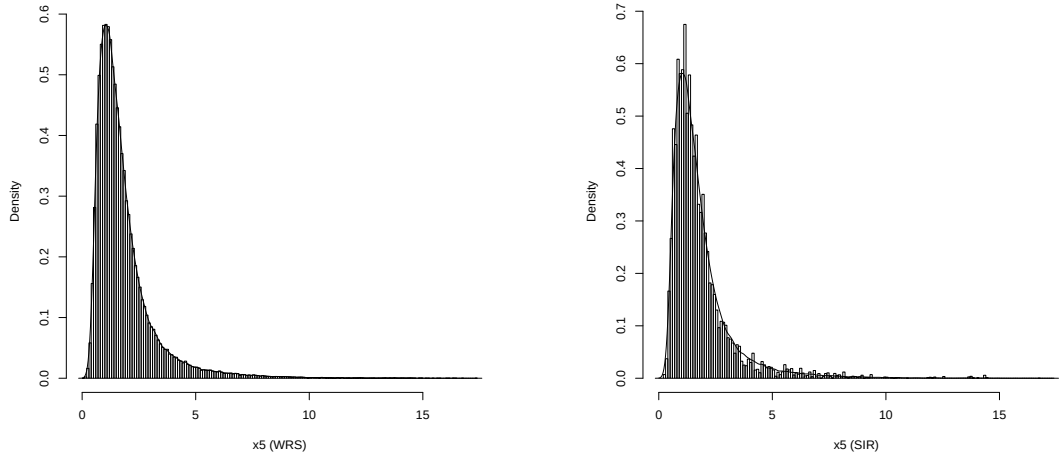

(a) WRS ($w = 4$)

(b) SIR ($\approx 0.06\%$ distinct)

Figure 11: Section 4.4 Example: Means for marginal components in $100,000$ draws from $\pi(x_{0:10}|z_{1:10})$ using perfect 11-dimensional rejection sampling, the SIR algorithm, and the WRS algorithm with $w = 9$.

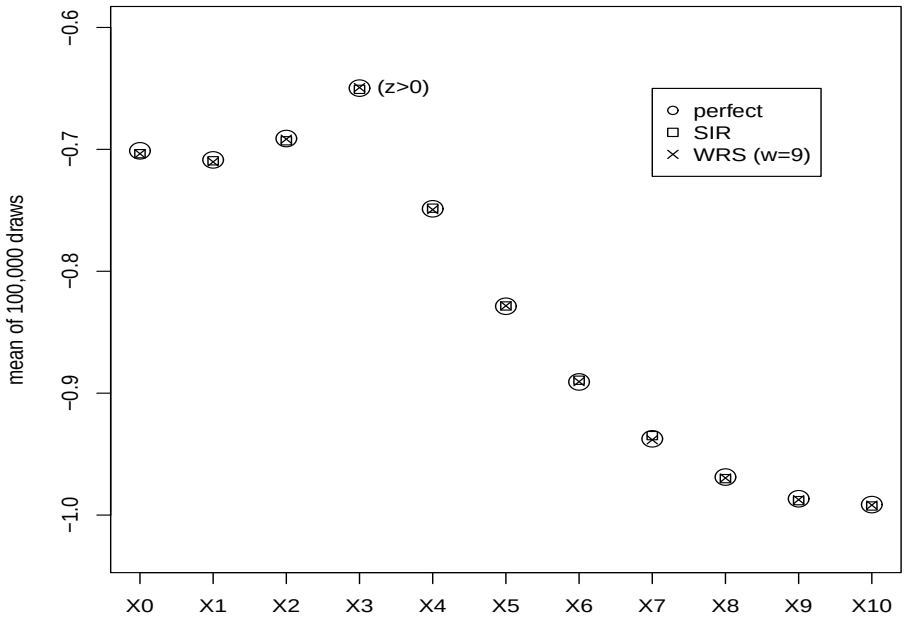

Figure 12: Section 4.4 Example: Means for certain marginal components in $100,000$ draws from $\pi(x_{0:1000}|z_{1:1000})$ using the SIR algorithm and the WRS algorithm with $w = 9$.

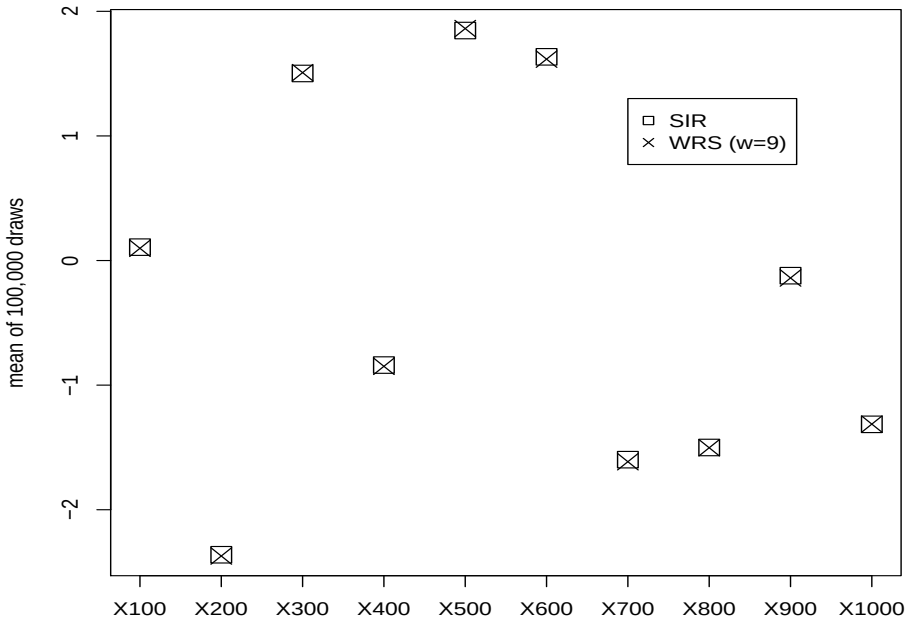

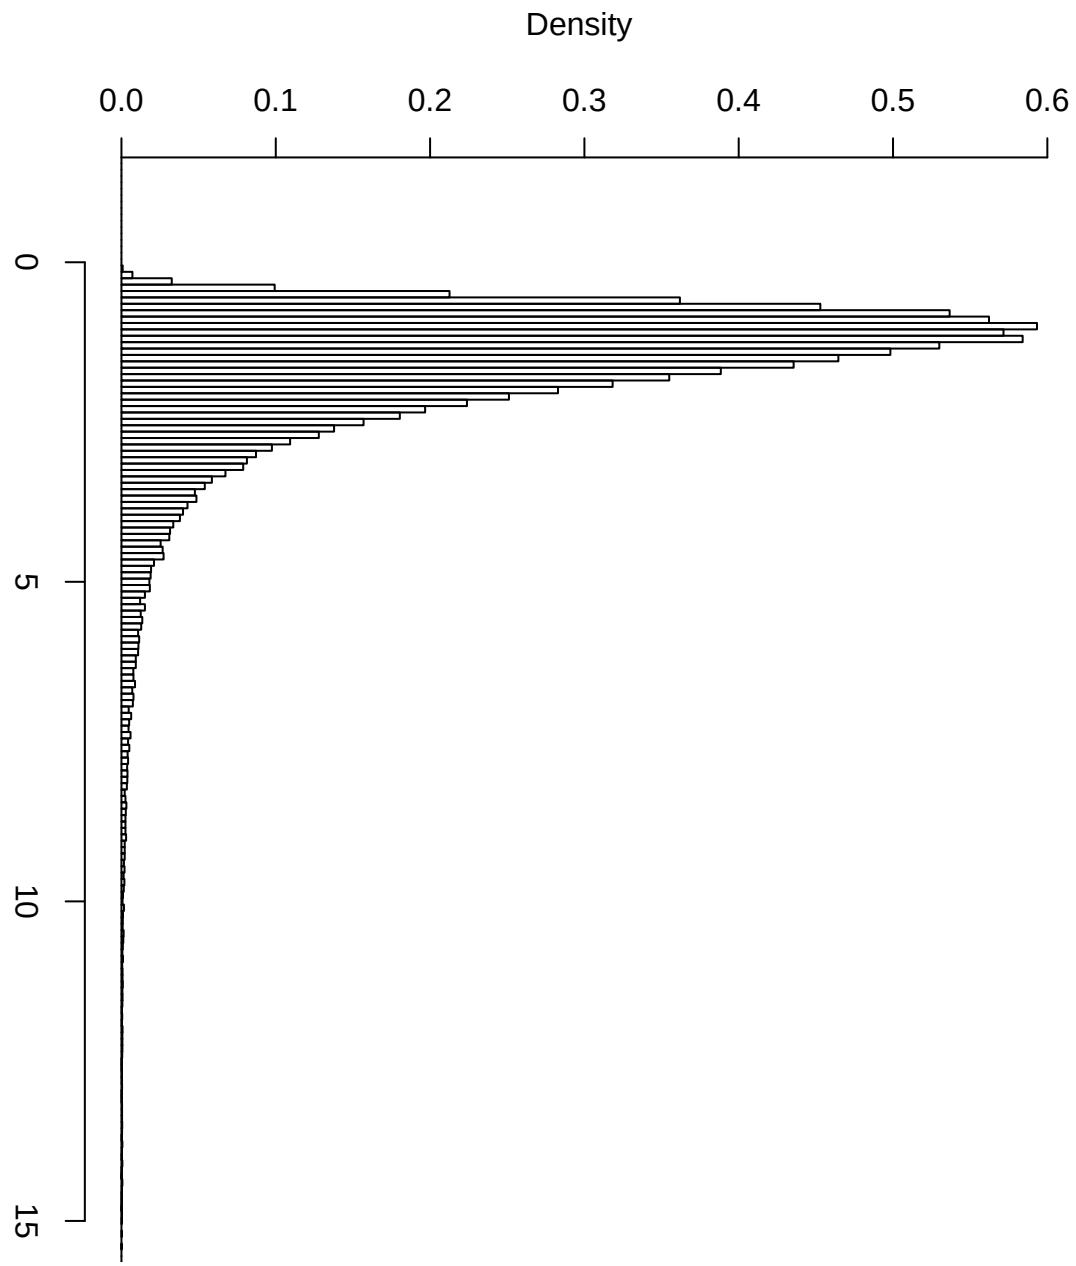
Density
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0
5
10
15

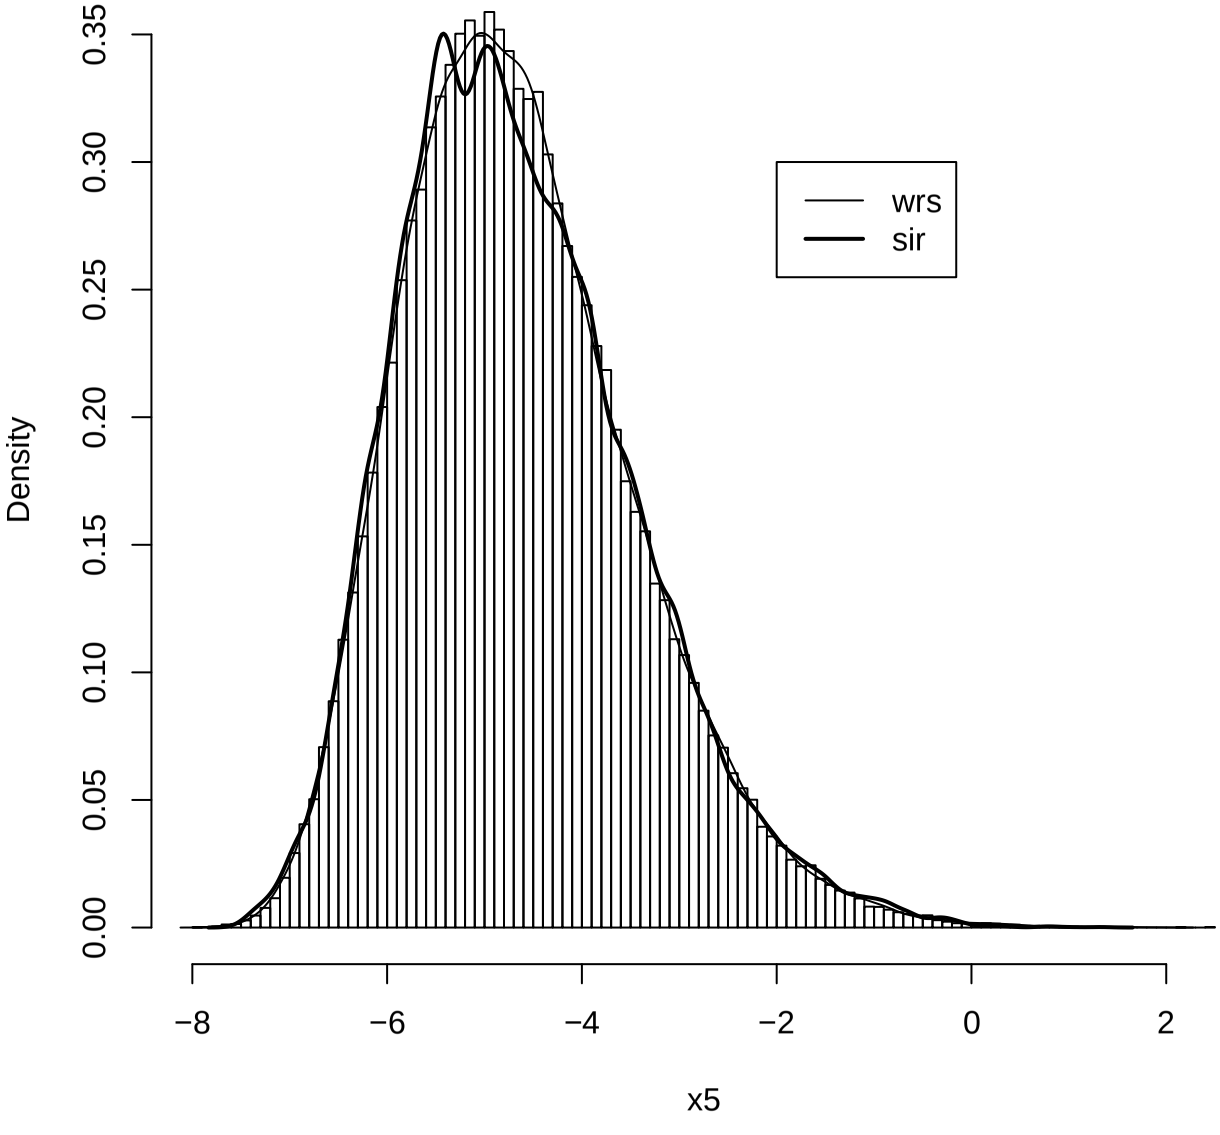
wrs
sir
Density
0.00
0.05
0.10
0.15
0.20
0.25
0.30
0.35
−8
−6
−4
−2
0
2
x5